



SMART 2020

The Ninth International Conference on Smart Systems, Devices and Technologies

ISBN: 978-1-61208-805-1

September 27th – October 1st, 2020

SMART 2020 Editors

Pascal Urien, Telecom Paris, France

Sergio Ilarri, University of Zaragoza, Spain

SMART 2020

Foreword

The Ninth International Conference on Smart Systems, Devices and Technologies (SMART 2020), held between September 27 – October 1st, 2020, continues a series of events covering tendencies towards future smart cities, specialized technologies and devices, environmental sensing, energy optimization, pollution control and socio-cultural aspects.

Digital societies take rapid developments toward smart environments. More and more social services are digitally available to citizens. The concept of ‘smart cities’ including all devices, services, technologies and applications associated with the concept sees a large adoption. Ubiquity and mobility added new dimensions to smart environments. Adoption of smartphones and digital finder maps, and increasing budgets for technical support of services to citizens settled a new behavioral paradigm of city inhabitants.

We take here the opportunity to warmly thank all the members of the SMART 2020 Technical Program Committee, as well as the numerous reviewers. The creation of such a broad and high quality conference program would not have been possible without their involvement. We also kindly thank all the authors who dedicated much of their time and efforts to contribute to SMART 2020. We truly believe that, thanks to all these efforts, the final conference program consisted of top quality contributions.

Also, this event could not have been a reality without the support of many individuals, organizations, and sponsors. We are grateful to the members of the SMART 2020 organizing committee for their help in handling the logistics and for their work to make this professional meeting a success.

We hope that SMART 2020 was a successful international forum for the exchange of ideas and results between academia and industry and for the promotion of progress in the field of smart systems, devices and technologies.

SMART 2020 Chairs:

SMART 2020 Publicity Chair

Mar Parra, Universitat Politecnica de Valencia, Spain

SMART 2020

COMMITTEE

SMART 2020 Publicity Chair

Mar Parra, Universitat Politecnica de Valencia, Spain

SMART 2020 Technical Program Committee

Lounis Adouane, Université de Technologie de Compiègne / Heudisayc, France
Ramakalyan Ayyagari, National Institute of Technology Tiruchirappalli, India
Karolina Baras, University of Madeira, Portugal
Lasse Berntzen, University of South-Eastern Norway, Norway
DeJiu Chen, KTH Royal Institute of Technology, Sweden
Patrizio Dazzi, ISTI-CNR, Italy
María del Mar Gómez Zamora, Universidad Rey Juan Carlos, Spain
Marta Sylvia del Río, Universidad de Monterrey, Mexico
Rachid El Bachtiri, USMBA University, Morocco
Steffen Fries, Siemens AG, Germany
Marco Furini, University of Modena and Reggio Emilia, Italy
Angelo Furno, LICIT laboratory - University of Gustave Eiffel / University of Lyon, France
Ivan Ganchev, University of Limerick, Ireland / University of Plovdiv "Paisii Hilendarski", Bulgaria
Maria Helena Gomes de Almeida Gonçalves Nadais, University of Aveiro, Portugal
Jorge J. Gomez-Sanz, Universidad Complutense de Madrid, Spain
Fakhrul Hazman Yusoff, Universiti Teknologi MARA, Malaysia
Gerold Hölzl, University of Passau, Germany
Tzung-Pei Hong, National University of Kaohsiung, Taiwan
Javier Ibanez-Guzman, Renault, France
Sergio Ilarri, University of Zaragoza, Spain
Simeon Keates, Edinburgh Napier University, UK
Abdelmajid Khelil, Landshut University of Applied Sciences, Germany
Pinar Kirci, Bursa Uludag University, Turkey
Kirill Krinkin, Saint-Petersburg Electrotechnical University, Russia
Wolfgang Leister, Norsk Regnesentral, Norway
Giovanni Livraga, Università degli Studi di Milano, Italy
Giuseppe Mangioni, University of Catania, Italy
Manuel Mastrofini, Università degli studi di Roma Tor Vergata & LAST Horizon, Italy
Natarajan Meghanathan, Jackson State University, USA
Dmitry Namiot, Lomonosov Moscow State University, Russia
Wilma Penzo, DISI - University of Bologna, Italy
Cathryn Peoples, Ulster University, UK
Christine Perakslis, Johnson & Wales University, USA
Ermanno Pietrosevoli, International Centre for Theoretical Physics, Italy
Marco Polignano, University of Bari Aldo Moro, Italy
Sherif Rashad, Morehead State University, USA
José M. Reyes Álamo, CUNY - New York City College of Technology, USA

Daniele Riboni, University of Cagliari, Italy
Michele Risi, University of Salerno, Italy
Marius Rohde Johannessen, University of South-Eastern Norway - School of Business, Norway
Enrique Romero-Cadaval, University of Extremadura, Spain
Demetrios Sampson, University of Piraeus, Greece
Farhan Siddiqui, Dickinson College, USA
Young-Joo Suh, POSTECH, Korea
Konstantinos Votis, Information Technologies Institute | Centre for Research and Technology Hellas, Thessaloniki, Greece
Qinggang Wu, Second Research Institute of Civil Aviation Administration of China, China
Mudasser F. Wyne, National University, USA
Wuyi Yue, Konan University, Japan
Sherali Zeadally, University of Kentucky, USA
Sotirios Ziavras, New Jersey Institute of Technology, USA

Copyright Information

For your reference, this is the text governing the copyright release for material published by IARIA.

The copyright release is a transfer of publication rights, which allows IARIA and its partners to drive the dissemination of the published material. This allows IARIA to give articles increased visibility via distribution, inclusion in libraries, and arrangements for submission to indexes.

I, the undersigned, declare that the article is original, and that I represent the authors of this article in the copyright release matters. If this work has been done as work-for-hire, I have obtained all necessary clearances to execute a copyright release. I hereby irrevocably transfer exclusive copyright for this material to IARIA. I give IARIA permission to reproduce the work in any media format such as, but not limited to, print, digital, or electronic. I give IARIA permission to distribute the materials without restriction to any institutions or individuals. I give IARIA permission to submit the work for inclusion in article repositories as IARIA sees fit.

I, the undersigned, declare that to the best of my knowledge, the article does not contain libelous or otherwise unlawful contents or invading the right of privacy or infringing on a proprietary right.

Following the copyright release, any circulated version of the article must bear the copyright notice and any header and footer information that IARIA applies to the published article.

IARIA grants royalty-free permission to the authors to disseminate the work, under the above provisions, for any academic, commercial, or industrial use. IARIA grants royalty-free permission to any individuals or institutions to make the article available electronically, online, or in print.

IARIA acknowledges that rights to any algorithm, process, procedure, apparatus, or articles of manufacture remain with the authors and their employers.

I, the undersigned, understand that IARIA will not be liable, in contract, tort (including, without limitation, negligence), pre-contract or other representations (other than fraudulent misrepresentations) or otherwise in connection with the publication of my work.

Exception to the above is made for work-for-hire performed while employed by the government. In that case, copyright to the material remains with the said government. The rightful owners (authors and government entity) grant unlimited and unrestricted permission to IARIA, IARIA's contractors, and IARIA's partners to further distribute the work.

Table of Contents

Blockchain Benefits for MAAS <i>Pascal Urien</i>	1
Simulation and Benchmarking of IoT Device Usage Scenarios Using Zephyr and Qemu <i>Bill Schirrmeister, Frank Geyer, and Steffen Spathe</i>	6

Blockchain Benefits for MAAS

Pascal Urien

Information Processing and Communication Laboratory (LTCI)
 Telecom Paris
 Saclay, France
 Pascal.Urien@Telecom-Paris.fr

Abstract—This paper introduces blockchain services for Mobility As A Service (MAAS) infrastructures, allowing to use public or private modes of transport during a trip. The blockchain system interacts with three classes of blockchain accounts: User, MAAS, and transport operators. It details some possible scenari and associated transactions. It demonstrates that this concept is realistic with emerging and existing technologies.

Keywords- Blockchain; Trust; Security; NFC

I. INTRODUCTION

The Mobility As A Service (MAAS, [1]) is an emerging concept, born in 2014 in Finland [2], allowing to use public or private modes of transport during a trip (see Figure 1). It involves the unification of mobility services, and a fusion of ticketing and multimodal information tools. According to [2] "The ecosystem consists of the transport infrastructure, transportation services, transport information and payment services".

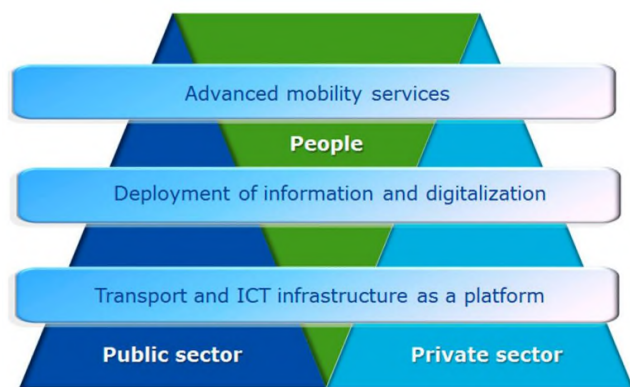


Figure 1. Mobility As A Service, MAAS

In this smart city context [10], many modes of transport are implemented such as bicycles, electric scooters, scooters, electric cars, buses, metro, and taxis. Some of them (generally private) rely on mobile applications associated with bank cards; others require tickets delivered by dedicated machines or subscription cards. In Helsinki, the WHIM app [9] offers several subscriptions; for example a monthly

package of 499 Euros, including unlimited public transport, taxis (up to 5 km), car-sharing and cycling.

In this paper we introduce the concept of blockchain services for MAAS. The MAAS is a federation of transport operators, managing ticketing services and validators. The main idea is to store these tickets in blockchain interacting with several entities (transport operator, user, MASS operators) as illustrated by Figure 2.

The paper is constructed according to the following outline. Section 2 presents blockchain benefits for MAAS, and some existing concepts. Section 3 introduces our use case. Section 4 details the functional architecture. Finally section 5 concludes this paper.




Figure 2. Actors involved in blockchain services for MAAS.

II. STATE OF ART

The first blockchain was deployed in 2009 by Satoshi Nakamoto [3], an anonymous creator (or group of creators). Bitcoin has had remarkable success, with a price reaching \$20,000 at the end of 2017, for a valuation of \$ 340 billion. Originally the blockchain carried out the transfer of crypto currency, but since 2015, several platforms appeared, such as Ethereum [4] or Hyperledger, which also allow the exchange of data or the execution of programs (smart contracts) charged in crypto currency.

The services delivered by a blockchain are as follows:

- Generation of signed transactions:
- Gathering of transactions by a P2P (Peer To Peer) network;

- Creation of transaction blocks by nodes of the P2P network.
- Chaining a new block to the existing block list (the blockchain), according to a consensus mechanism such as PoW (*Proof of Work*) or PoS (*Proof of Stake*).

The main contribution of blockchain is the creation of trust, based on consensus, i.e., there is no trusted third party. Transactions are authenticated by cryptographic signatures and are time-stamped, when they are mined (i.e., inserted in a block). The database storing these blocks is very widely duplicated in nodes of the P2P network; WEB APIs allow ubiquitous and secure access to data hosted by the blockchain infrastructure.

In the absence of a trusted third party, participants generate their keys. A private key is usually a 32 bytes (256 bits) random number, from which a public key and the blockchain account identifier are calculated. Blockchain is by nature based on a consensus of players, and therefore could be the technological cornerstone of the new MAAS mobility, bringing together multiple players.

An example of this concept is the "planar networks" project [5], which uses Ethereum smart contracts (see Figure 3), for implementation of universal ticketing system.




Figure 3. The planar networks project uses Smart Contracts for implementation of universal ticketing system.

The creation of crypto currencies dedicated to MAAS is also an opportunity. The DOVU project [6] deploys a crypto currency (see Figure 4) dedicated to multimodal mobility services.




Figure 4. The DOVU project

III. USE CASE

This section details an experimental implementation of the concept of "Blockchain for MAAS". The latter uses the ROPSTEN (Ethereum) test blockchain, a Near Field


Communication (NFC) chip card performing transaction signature, and free software (BTOOLS [7]).

A. Blockchain Ticket



Figure 5. A French SNCF ticket

In order to clarify our approach, we consider the e-ticket illustrated in Figure 5. Reading the QR-Code (located at the top right) by a dedicated validator allows access to the mobility service. The ticket (content of the QR code) is a series of ASCII characters illustrated by the left part of Figure 6.



@User = 1f889b79e2e1d4c412b2eea5e5286e44884414b5

Figure 6. Legacy Ticket (left) and blockchain ticket (right)

B. Blockchain MASS entities

Three types of blockchain accounts are involved in MAAS:

- User accounts, using addresses, deduced from previously generated private keys.
- MAAS entities, identified by their addresses.
- Transport operators, identified by their addresses.

Figure 7 lists the MAAS actors (MASS entity, user, transport operator), and their associated addresses in the ROPSTEN blockchain. Transactions can be retrieved at [11].

Entity	Ropsten Blockchain Address
User account	1f889b79e2e1d4c412b2eea5e5286e44884414b5
MAAS	266abe4b01ce1899100c36c3d33c872b52cb12d8
Transport operator	11baa7742ff512acc35b86d2c1b02391d533a4b2

Figure 7. MAAS actors and their ROPSTEN addresses.

IV. FUNCTIONAL ARCHITECTURE

The functional architecture of MAAS is illustrated in Figure 8. Three types of services are supported, the ticket generation, the ticket transfer to user, and the ticket use in a transport network.

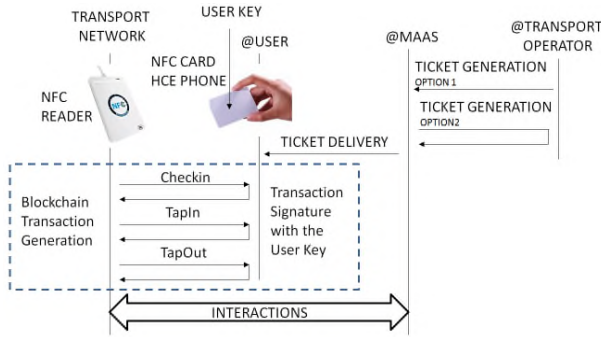


Figure 8. Functionnal Architecture

A. Ticket Generation

Transaction Hash:	0x8069bda915d10b60da37983c234497438166b005356633192fc24b24f741a567
Status:	Success
Block:	5042109 3013000 Block Confirmations
Timestamp:	477 days 31 mins ago (Feb-18-2019 08:56:05 AM +UTC)
From:	0x11baa7742ff512acc35b86d2c1b02391d533a4b2 Transport Operator
To:	0x266abe4b01ce1899100c36c3d33e872b52cb12d8 MASS
Value:	0 Ether (\$0.00)
Transaction Fee:	0.00088092 Ether (\$0.000000)
Gas Limit:	60,000
Gas Used by Transaction:	29,364 (48.94%)
Gas Price:	0.00000003 Ether (30 Gwei)
Nonce	0 1
Input Data:	0xi0CV UGDQ*W 345654746 1211 FRP*YOFRRNS0 8009 11/02 0029090162200235586 1f889b79e2e1d4c412b2eea5e5286e44884414b5 2FA000 00000

Figure 9. Ticket generation from the transport operator to the MASS.

The basic idea is to store the ticket in the blockchain account of the MAAS entity. Two modes of ticket generation are possible:

- Option 1: the transport operator generates the ticket in a transaction, whose recipient is the MAAS entity. This option is illustrated in Figure 9. This transaction can be retrieved at [12].
- Option 2: MAAS generates the ticket in a self addressed transaction. This transaction is illustrated in Figure 10; it can be retrieved from [13].

Transaction Hash:	0x64b915bab898c38953e488a10e9867ffaba35b933e7f64cb508a200146b41ed5
Status:	Success
Block:	5039041 3016706 Block Confirmations
Timestamp:	477 days 11 hrs ago (Feb-17-2019 09:39:48 PM +UTC)
From:	0x266abe4b01ce1899100c36c3d33e872b52cb12d8 MASS
To:	0x266abe4b01ce1899100c36c3d33e872b52cb12d8 MASS
Value:	0 Ether (\$0.00)
Transaction Fee:	0.00088092 Ether (\$0.000000)
Gas Limit:	60,000
Gas Used by Transaction:	29,364 (48.94%)
Gas Price:	0.00000003 Ether (30 Gwei)
Nonce	2 0
Input Data:	0xi0CV UGDQ*W 345654746 1211 FRP*YOFRRNS0 8009 11/02 0029090162200235586 1f889b79e2e1d4c412b2eea5e5286e44884414b5 2FA000 00000

Figure 10. Ticket generation from the MASS to MASS.

B. Ticket Transfer to User

The MAAS transfers the ticket to the client in a transaction, as illustrated in Figure 11. The ticket is thereafter stored in the client's blockchain account. This transaction can be retrieved at [14].

Transaction Hash:	0xaabaed08f487896f6d1aed70b01da025f06d0feb626b60f58068bed2fc1dbc
Status:	Success
Block:	5038990 3016884 Block Confirmations
Timestamp:	477 days 12 hrs ago (Feb-17-2019 09:28:20 PM +UTC)
From:	0x266abe4b01ce1899100c36c3d33e872b52cb12d8 MASS
To:	0x1f889b79e2e1d4c412b2eea5e5286e44884414b5 User
Value:	0.01 Ether (\$0.00)
Transaction Fee:	0.00087888 Ether (\$0.000000)
Gas Limit:	60,000
Gas Used by Transaction:	29,296 (48.83%)
Gas Price:	0.00000003 Ether (30 Gwei)
Nonce	1 2
Input Data:	0xi0CV UGDQ*W 345654746 1211 FRP*YOFRRNS0 8009 11/02 0029090162200235586 f889b79e2e1d4c412b2eea5e5286e44884414b5 2FA000 00000

Figure 11. Ticket Transfer from MAAS to user

C. Ticket Use

The user is equipped with a Near Field Communication (NFC) card or a Host Card Emulation (HCE) smartphone. The NFC card (possibly virtual) hosts the private key associated to the blockchain account, and is able of carrying out the transaction signature.




Figure 12. Illustration of the *CheckIn* scenario.

In an off-line approach (*Checkin* scenario, see Figure 12), the NFC card stores a ticket linked to an address, whose authenticity is proven by a signature, calculated using the private key. The user's blockchain address can be associated with a certificate issued by the MAAS, if a static authentication mechanism (off-line) is required.

In an on-line approach (*TapIn, TapOut* scenario), similar to payment by contactless bank card in the London Underground, proof of customer identity is required at the start and end of the mobility service. The ticket (or contract) can be stored in the MAAS account, or paid by using a dedicated crypto currency.



Figure 13. Illustration of the *TapIn, TapOut* scenario

1) *Checkin* Transaction.

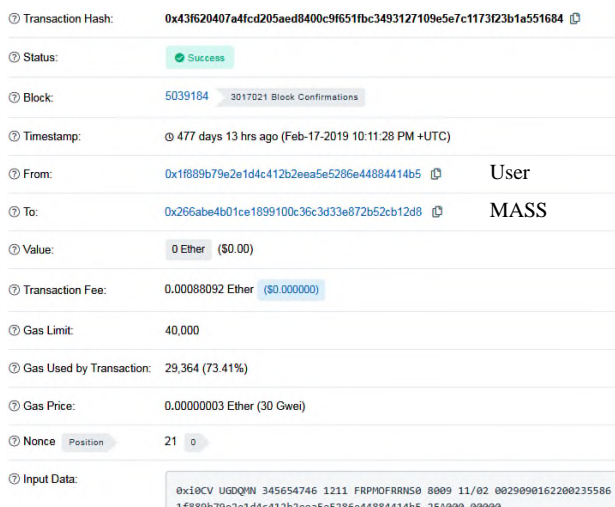


Figure 14. Illustration of a *Checkin* Transaction

Figure 14 illustrates a transaction between user and MAAS entity which contains the ticket. It can be retrieved at

[15]. This latter is authenticated by the signature generated using the private key bound to the user account.

2) *TapIn, TapOut* Transactions.

Figures 15 and 16 illustrate respectively the *TapIn* and *TapOut* procedures. These transactions can be respectively retrieved at [15] and [16]. A self transaction, addressed to the user, proves its identity, and is thereafter time stamped by the blockchain infrastructure

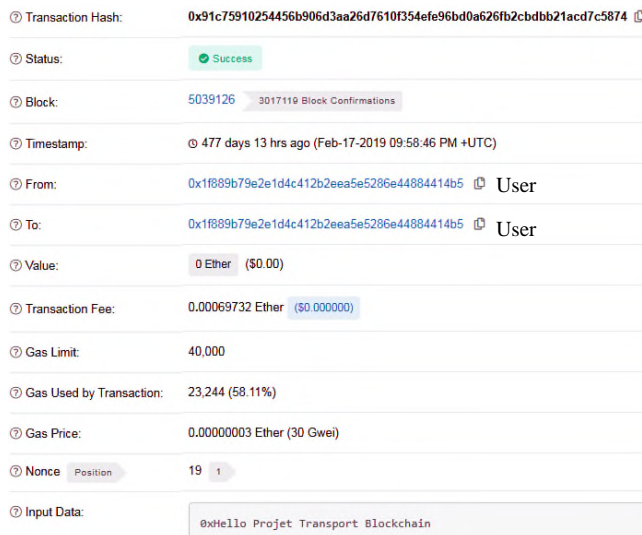


Figure 15. A *TapIn* transaction, proving the user identity, recorded and time stamped by the blockchain.

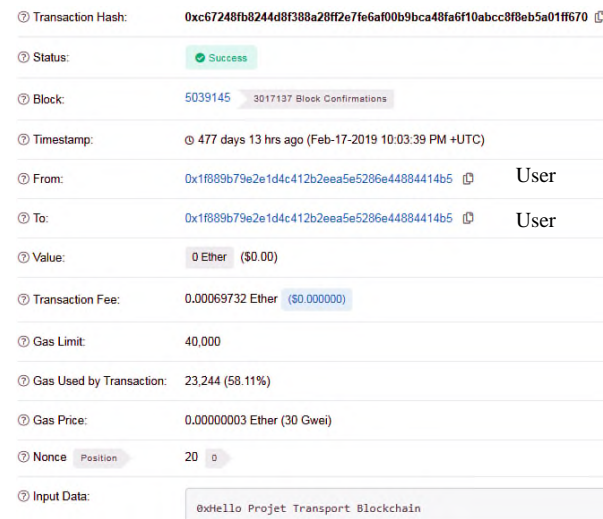


Figure 16. A *TapOut* transaction, proving the user identity, recorded and time stamped by the blockchain.

V. CONCLUSION

In this paper we presented a blockchain infrastructure dedicated to MAAS services, which comprises three actors, MAAS entity, transport operator and user. A ticket can be generated in MAAS to MASS transaction or in transport

operator to MAAS transaction. The MAAS transfers this ticket to user account. Validators of the transportation system check ticket according to two operating modes *CheckIn* and *TapIn-TapOut*. In the first mode the user generates a transaction to the MAAS entity in which the ticket is inserted; the ticket is bound to user's identity and its use is time stamped by the blockchain. In the second mode the user proves its blockchain identity by self-addressed transaction; thereafter the transport system checks its subscription.

In summary blockchain could create very attractive technological platform for MAAS services, due to the lack of third trusted party. Nevertheless such infrastructures should provide the required level of performances.

REFERENCES

- [1] https://en.wikipedia.org/wiki/Mobility_as_a_service, accessed September 2020
- [2] "Mobility as a Service - the new transport paradigm". June 2014, <https://www.lvm.fi/>, accessed September 2020
- [3] S. Nakamoto, S., "Bitcoin: A peer-to-peer electronic cash system", 2008
- [4] W. Gavin, Ethereum Yellow Paper, "Ethereum : a Secure Decentralised Generalised Transaction Ledger", EIP-150, 2015, <http://yellowpaper.io>, seen september 2020.
- [5] The Planar Network Blockchain Transport, <https://planar.network/assets/docs/whitepaper.pdf>, retrieved September 2020.
- [6] DOVU: a unified token, wallet and marketplace for earning and spending mobility related rewards, <https://dovu.io/>, retrieved September 2020.
- [7] P. Urien et al., "BTOOLS: Trusted Transaction Generation for Bitcoin and Ethereum Blockchain Based on Crypto Currency SmartCard", ICDT 2018, The Thirteenth International Conference on Digital Telecommunications
- [8] P. Urien, "Towards secure elements for trusted transactions in blockchain and blockchain IoT (BIoT) Platforms", Fourth International Conference on Mobile and Secure Services (MobiSecServ), 2018.
- [9] WHIM application, <https://whimapp.com/>, retrieved September 2020.
- [10] A. Zanella, N. Bui, A. Castellani, L. Vangelista and M. Zorzi, "Internet of Things for Smart Cities," in IEEE Internet of Things Journal, vol. 1, no. 1, pp. 22-32, Feb. 2014.
- [11] Ropsten Testnet Explorer, <https://ropsten.etherscan.io>, retrieved September 2020.
- [12] Ticket generation from the transport operator to the MASS. <https://ropsten.etherscan.io/tx/0x8069bda915d10b60da37983c234497438166b005356633192fc24b24f741a567>, retrieved September 2020
- [13] Ticket generation from the MASS to MASS, <https://ropsten.etherscan.io/tx/0x64b915bab898c38953e488a10e9867ffaba35b933e7f64cb508a200146b41ed5>, retrieved September 2020
- [14] Ticket Transfert from MAAS to user, <https://ropsten.etherscan.io/tx/0xaabaed08f4f87896f6d1aed70b01da025f06d0feba62eb60f58068bed2fc1dbc>, retrieved September 2020.
- [15] Checkin Transaction, <https://ropsten.etherscan.io/tx/0x43f620407a4fcd205aed8400c9f651fbc3493127109e5e7c1173f23b1a551684>, retrieved September 2020
- [16] TapIn transaction, <https://ropsten.etherscan.io/tx/0x91c75910254456b906d3aa26>

d7610f354efe96bd0a626fb2cbdbb21acd7c5874, retrieved September 2020

- [17] TapOut transaction, <https://ropsten.etherscan.io/tx/0xc67248fb8244d8f388a28ff2e7fe6af00b9bca48fa6f10abcc8f8eb5a01ff670>, retrieved September 2020

Simulation and Benchmarking of IoT Device Usage Scenarios Using Zephyr and Qemu

Bill Schirrmeister, Frank Geyer, Steffen Späthe

Friedrich Schiller University Jena
Department of Computer Science, Software Engineering Group
Jena, Germany

Email: bill.schirrmeister@uni-jena.de
frank.geyer@uni-jena.de
steffen.spaethe@uni-jena.de

Abstract—The development of a device command and control infrastructure for Internet of Things devices with a focus on low resource consumption is the primary goal of the research project "unic²ast". This paper proposes a simulation environment that enables to carry out benchmarks of this infrastructure according to a variety of application scenarios. The presented simulation setup uses Qemu virtualized embedded devices with an application based on Zephyr OS. The basic suitability of the approach is demonstrated and potentials for further development are identified.

Keywords—IoT Device Management; Lightweight M2M; Qemu; Zephyr; Device Simulation.

I. INTRODUCTION

The number of connected devices in the Internet of Things (IoT) is forecast to reach over 75 billion by 2025 [1]. As the overall goal of the research project unic²ast [2] a device command and control infrastructure (DCCI) for scenarios with a large number of IoT devices is meant to be realized.

Following the well-known client-server-model, a server infrastructure for managing and administrating IoT devices which provides transparency with regard to their concrete application scenarios is being developed. A good system performance is considered to be essential for this project, since it is required to manage and integrate a potentially large amount of devices (a) with possibly constraint hardware resources (b). To evaluate this system property, development of a simulation system is one of the project goals. This should allow for the measuring of values that are as realistic as possible and help to identify bottlenecks in the processing inside the DCCI system.

Available network simulation systems such as "ns-3" [3] or "OMNeT++" [4] are not sufficient for unic²ast, since models of the server and the processing logic would need to be created. This would lead to a rather difficult realization and cannot be achieved without knowledge about the runtime behavior of the individual components. The term "individual components" also includes the operating system, hardware (input / output to RAM, hard disk, etc.), runtime environment (e.g., Java VM), database management systems (MariaDB, Postgres, etc.), and other web services. Likewise, unic²ast has to use specific protocols (such as OMA Lightweight M2M (LwM2M)) in order to be able to provide the desired system functionalities.

Existing network simulation systems do not offer out-of-the-box implementation of such specific IoT device protocols.

Unic²ast relies on the use of a real DCCI instance that is occupied by virtual IoT devices (i.e., simulated load) in order to test the remote maintenance of multiple IoT devices via the DCCI with a focus on performance and stability analysis. The simulation is coordinated by a test coordinator, i.e., a software solution for planning, controlling and evaluating the simulation runs. Such a system structure is not uncommon in practice. Standard load test systems for server applications such as Apache JMeter [5] or Gatling [6] are built according to the same principle.

Unic²ast follows an approach where virtual IoT devices are used to test the real DCCI server implementation. In principle, one could roughly calculate load limits or determine them by using purely virtual simulations. However, it was opted for the possibly more resource-intensive evaluation approach because the server system must be able to handle a large number of any IoT devices in productive operation. Therefore, it is not possible to predict the need for communication in the form of static information, as required by other simulation approaches. In particular, bottlenecks in internal communication processes inside the device command and control infrastructure are unknown in advance of the load tests.

This paper is not intended to provide a comprehensive test specification for the DCCI. Likewise, there is no complete specification of the system's properties to be tested and no detailed definition of test procedures.

The aim of this paper is to present the infrastructural approach for the parallel execution of several, configurable client instances for a given application scenario with focus on benchmarks of the overall system. On the basis of the explanations here, it should be possible at further steps to define concrete and more complex test sequences and to carry out corresponding data collection with focus on performance-benchmarks at runtime.

We have organized the rest of this paper in the following four sections. The core requirements to the benchmark infrastructure and available implementation approaches are presented in Section II. Section III is about how the selected infrastructure approach was implemented in detail. Aspects of

a practical usage are shown in Section IV. At the conclusion of the paper, in Section V we summarize the results achieved and briefly evaluate the chosen approach based on our formulated requirements.

II. REQUIREMENTS AND IMPLEMENTATION APPROACHES

To prepare the selection of a suitable approach and environment for simulation within unic²ast, the following essential requirements to the benchmark infrastructure were defined:

- simulated devices behave quite similar to original devices
- ability to run multiple simulated device instances on the same host
- LwM2M connection to the system under test
- parameterization of the device instances
- retrieval of runtime information after startup
- automation of the device instance lifetime and the test procedure

OMA Lightweight M2M (LwM2M) in this context is a protocol of Open Mobile Alliance (OMA) for IoT device management [7]. The LwM2M protocol defines an application layer communication protocol between client and server and focuses on low resource consumption. LwM2M is based on Constrained Application Protocol (CoAP).

As described before, the benchmark infrastructure should use simulated client devices. A valid approach would be to implement the client's functionality within a dedicated simulation program. The program could be implemented in a comfortable high-level language, and all available libraries could be used. Furthermore, this program could use multithreading to simulate several individual devices simultaneously. This simulation program could, therefore, be implemented with relatively little effort. This approach would also lead to a rather simple test setup.

On the other hand, an application program already exists for the target hardware, which realizes the communication between server and embedded system. This embedded program itself is its best simulation in the context of a load test of the overall system. However, the use of several hundred real devices is inappropriate. Therefore, another approach to build a test system setup is to emulate the embedded devices, including their existing application logic. In this way, the additional parallel development of a synthetic load driver based on an additional development environment can be avoided.

III. APPROACH

A combination of shell scripts with "Zephyr OS" (Zephyr) and "Quick Emulator" (Qemu) was used to implement a solution that takes the given requirements into consideration. Zephyr OS is an open source real-time operating system for IoT devices with a small memory size and fixed hardware configuration [8]. Qemu is a free virtualization software for complete hardware emulation [9].

Figure 1 shows the resulting system landscape. The system to be tested is shown inside the *System under test* box as *DCCI*. However, the actual structure of the system to test is considered to be a black box and therefore was omitted. The landscape to be tested can consist of complex server structures.

The virtual system part (box *Simulated devices*) usually consists of several IoT devices that are implemented as Qemu processes executed in parallel. Each Qemu process simulates exactly one IoT device that is operated by a Zephyr OS based embedded application. For communication purposes between virtual devices and the system under test via LwM2M, an extended version of the LwM2M client included in Zephyr is used. The underlying Internet Protocol (IP) connection of clients is realized by setting up a virtual bridge connecting them with the hosts physical ethernet adapter.

Every virtual device is assigned one named pipe for input and output which is set up in the hosts file system. Their coupling with the respective Qemu process is realized by configuring it as a serial port when calling the Qemu executable. On the Zephyr OS side, this port is connected to the built-in shell subsystem. This makes it possible to not only send shell commands to each device, but also to record the shell output.

Functions for automating preparation, execution and post-processing of benchmarks are implemented in Bash scripts. A higher-level coordinator script defines the test structure and sequence and therefore uses the functions modularized in other scripts. The higher-level script thus represents the implementation of a concrete test scenario. Other scenarios can be realized by alternative implementations of this script.

In the following part the most important functions, based on prior defined requirements, will be discussed in more detail.

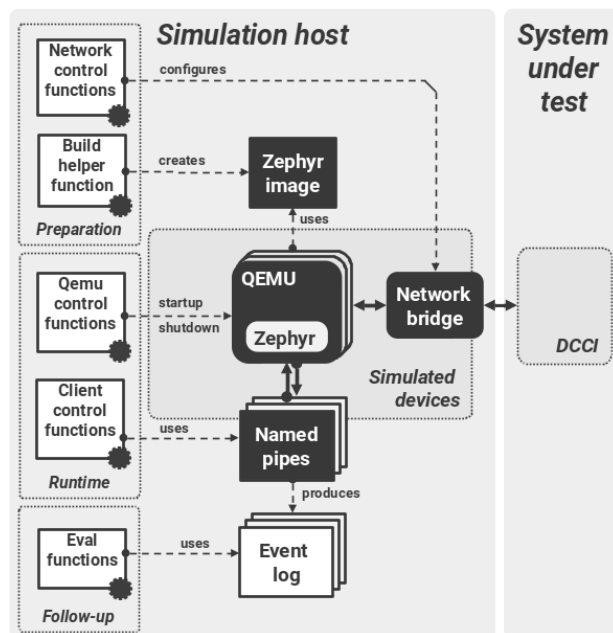


Figure 1. Simulation-landscape at unic²ast

A. Parallel Execution on a Single Host

Parallel execution of multiple test clients on a single host is done by launching several virtual machines at once.

The unic²ast project makes use of the `qemu_x86` board configuration shipped with Zephyr source code. Following the Zephyr documentation, the application program that implements the LwM2M client is developed based on the Zephyr kernel. It gets compiled into a single binary file alongside the

kernel to be then executed in a virtual machine. Execution is thus isolated from the host and other virtual machine instances. The compilation is done by calling `cmake` with the parameter

```
-DBOARD=qemu_x86
```

followed by an additional call to `make`. After successful compilation, the resulting binary file is then used when launching a Qemu process. This may be done by calling the respective Qemu executable via

```
qemu-system-i386 -kernel [PATH-TO-IMAGE]
```

providing it with the built Zephyr based application image via parameter `kernel` (additional parameters omitted due to layout restrictions).

By calling Qemu multiple times successively, several Qemu processes can be launched in parallel, each executing a client in its own virtual machine. To avoid the need to execute each call to Qemu individually on the command line, this was automated within a parameterizable loop as part of a Bash function.

A call to the specified function results in the following behavior

- creation of a named pipe per client
- generation of a unique medium access control (MAC) address per client
- startup of the specified number of Qemu processes, utilizing the built Zephyr based application image as described above, and assigning the named pipe to the serial port and the MAC address to the ethernet adapter respectively
- calling of `cat` on the output pipe for each client and redirecting its output stream to a file
- storage of the meta data describing the clients and other relevant data during runtime in corresponding files

Similarly, a shutdown function takes care of the scheduled shutdown of the started instances using the process ID noted as part of the meta data for each Qemu process.

B. Connection to Target System

The client's ability to connect to the target system using the IP is a necessary prerequisite for the required LwM2M connection, which is based upon CoAP and is achieved through the following four steps:

- creating a network bridge on the host
- providing an ethernet adapter to the Zephyr image
- linking the Qemu machine's ethernet adapters to the bridge on startup
- allocating a suitable unique IP address to each client

The possibility of setting up a network bridge as a virtual connection element on layer 2 of the Open Systems Interconnection (OSI) model is already built into modern Linux operating systems. The necessary bridge can therefore be set up in a terminal via

```
ip link add [BRIDGE] type bridge
```

After assigning an IP address to the host to ensure its reachability on the bridge, the host's physical ethernet adapter is also connected to the bridge via

```
ip link set [ADAPTER] master [BRIDGE]
```

The bridge thus forms a virtual extension of the physical network on the MAC layer that is executed on the host.

To successfully connect the clients to the bridge, they need an ethernet adapter too. According to Zephyr OS documentation for building `qemu_x86`, the source code required for the virtual ethernet adapter E1000 can be included in the compiled image by specifying the overlay file `overlay-e1000.conf` alongside the others with the parameter `-DCONF_FILE` when `cmake` is called.

The Qemu machines are then connected to the bridge by specifying

```
-nic bridge,br=[BRIDGE],mac=[MAC-ADDR]
```

during the call to the Qemu executable, where `MAC-ADDR` is the generated unique MAC address of a client.

In addition to the connection already described on the MAC layer, the clients still need a unique Internet Protocol version 4 (IPv4) address in order to be able to establish CoAP based connections to the target system. A static assignment as in the LwM2M sample found in Zephyr's sources is unfavorable because the address would need to be configured at compile time. If then arbitrary `n` instances of the client were desired to be assigned an address, potentially expensive `n` compilation runs would be necessary. The dynamic assignment of addresses via Dynamic Host Configuration Protocol (DHCP) has also proven to be impractical. With both a local DHCP server on the host and an external server on the network, it was impractical to reliably serve IPv4 addresses to a large number of volatile clients (in the hundreds).

Our solution to this problem uses static address assignment by means of a Zephyr shell command combined with address generation using `prips`. The command line program `prips` prints a line-by-line listing of all IPv4 addresses of a given IPv4 subnet. This subnet has to be specified in Classless Inter-Domain Routing (CIDR) notation when calling `prips` as follows

```
prips "192.168.0.0/24"
```

Subtracting at least the subnet's first address (network) and the last one (broadcast) is necessary afterwards, since they must not be used for client assignment. Further addresses may have to be omitted if they are already assigned to the client's host or other computers on the physical network, resulting in the final list of usable addresses. The assignment to the clients is then carried out by a dedicated Zephyr shell command, which was implemented specifically for that purpose. It then assigns the address inside its callback by calling the function

```
net_if_ipv4_addr_add_by_index
```

which is provided by Zephyr's network stack accordingly.

C. Parameterization and Value Retrieval at Runtime

A named pipe each for input and output streams in the host's file system are used for direct communication between

the host and the virtual machines during runtime. Splitting into two file paths is necessary because buffer conflicts could occur if only a single path is used for both input and output. The separate pipes have the same base file name, but differ in their extension, being assigned an `.in` or `.out`. The assignment of such a "double pipe" to its dedicated virtual machine then is made by specifying just the base file name as parameter value for the serial port when calling the Qemu executable. The host operating system automatically separates the input and output character streams to the separate files accordingly.

The parameterization of a client at runtime is then carried out by the transmission of specific shell commands to each Qemu machine. This may be achieved by using `echo` like

```
$ echo "[COMMAND]" > [PATH]/[PIPE].in
```

redirecting it to the respective pipe's input path. In return, the pipe's output path enables all shell outputs of a client to be recorded, for example by creating a `cat` process of the form

```
$ cat [PATH]/[PIPE].out >> [PATH]/[FILE] &
```

Since the basic shell commands available in Zephyr OS are insufficient for the purposes described here, application-specific commands had to be implemented.

The shell interface of a Zephyr application follows a hierarchical structure. Several subordinate command words can be assigned to a higher-level command word to form a specific command. Parameter values that have to be transmitted alongside a command can be specified on the leaves of the defined tree. The callback method, which is also to be specified on the leaves, is called when the respective command is received, providing it with the transferred parameter values and thus allows the behavior of the command to be implemented within the Zephyr application. Among other things, registering a command on the Zephyr shell subsystem can be done through special C macros [10].

The concrete commands for benchmarking implemented in this concept allow the necessary configuration of parameters that are required for the secure connection to the target system (e.g., Datagram Transport Layer Security (DTLS) pre-shared key). In addition to that, virtual temperature sensors may also be created on each client, the values of which are randomly varied at random time intervals in order to generate a realistic load by means of subscriptions made by the server.

Input or output via named pipes with the above shell commands can be issued manually on the host's shell. However, these processes were largely automated by script functions.

D. Client Identities and Resource Subscriptions

For securing the client's CoAP/LwM2M connection to the target system by means of the DTLS protocol within the context of the `uni2ast` project, the pre-shared key approach was to be used.

This key which has to be unique for each client must be known to the server in advance alongside its unique identifier string and the LwM2M endpoint name. In order to carry out load tests of the target system with varying multitudes of these clients information stored to the target's database. Therefore, it is necessary to provide functions that enable this data

- to be reported to the target system at the start of each individual test run and
- to be deleted from the target system after their end, respectively.

For the same reason and because of the number of clients in a test run being potentially large, it is also required to automatically generate the client's identity and security information parts, as mentioned above.

Randomly generating these values is undertaken in the implementing Bash function by read access to `/dev/urandom` of the Linux host. They are collected in a Comma Separated Value (CSV) file and are thus available to the other implemented Bash functions that need to have access to them, for example the Bash functions explained in Subsection III-C that are used for parameterizing the clients. Furthermore, the functions responsible for reporting the clients information to the target and deleting them later are implemented to use a Representational State Transfer Application Programming Interface (REST API) provided by the target specifically for that purpose.

IV. PRACTICAL USAGE

To demonstrate the presented solution, a simple test scenario defined for the real DCCI system implementation developed in `uni2ast` and the experiences made when preparing and performing the benchmark are given below.

A. Example Test Scenario Definition

The objective of the chosen scenario was to allow for the measurement of the target system's response times as seen by the clients, as well as the resource requirements of the target system in the context of an increasing number of clients from test run to test run. Based on the assumption that the regular registration update sequences between each client device and the DCCI server and the notification messages sent by the clients due to server side subscriptions for changing values are sufficient to cause a certain load on the target system, the scenario was defined to provide that

- each client simulates two temperature sensors (instances of the Internet Protocol for Smart Objects (IPSO) object 3303) with values adjusted at random intervals and
- the server subscribes to changes of these values as soon as the clients are connected.

Data on the target system's performance should then be gathered by measuring the registration event response latency of the server on client side and the server's computing and working memory utilization on server side during a fixed time interval the coordinator script has to wait during each run of the test. The procedure in this scenario should consist of four consecutive runs, starting with 100 clients during the first run, increasing to 400 and 800 throughout the second and third runs and ending up with 1,000 clients during the last run. The waiting period before shutting down the clients at the end of each run was set to 90 seconds.

B. Coordinator Implementation

The scenario described was implemented as a coordinator script, according to the principle explained in Section III, and uses the Bash functions implemented as part of the solution. The practical process flow for test setups is shown in Figure 2.

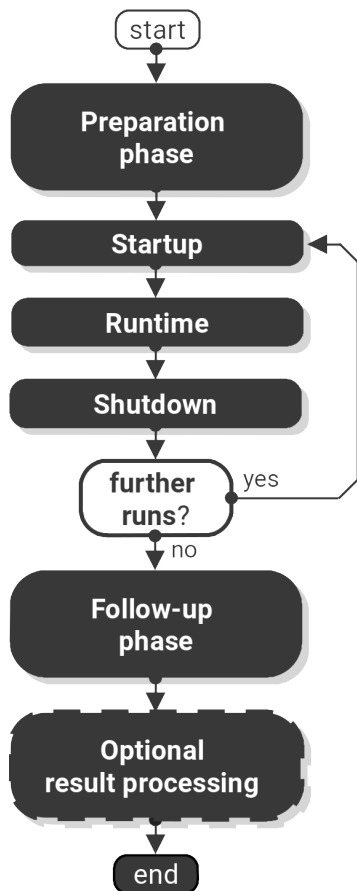


Figure 2. process flow in unic²ast test setup

The Preparation phase includes tasks that are necessary for the test execution, but are not to be repeated as part of the test in each run. This applies, for example, to build the Zephyr-based application image for the virtualized clients, creating the network bridge on the host but also booting up the target system and any other necessary steps to ensure that the target system can be reached from the test host.

Similarly, the Follow-Up phase includes all activities that are carried out independently of the individual test runs after the completed test. This includes for example removing the network bridge, possibly shutting down the target system, and evaluating the log files generated on the target system and on the clients by using R.

The three main phases, however, depend on the design of the specific test scenario.

The Startup phase includes at least:

- generating the security information of the clients and announce them to the target system
- starting the clients and assigning the respective IP addresses

- configuring the clients with the respective security information and the simulated temperature sensors

The following Runtime phase includes at least:

- connecting the clients to the server via LwM2M registration
- setting up the subscriptions by the server
- waiting for the period of time specified before the start of the test
- disconnecting all clients via LwM2M deregistration

The end of a test run with the Shutdown phase is indicated by:

- terminating the Qemu processes
- deleting the client’s security information from the server
- collecting the log files of all clients and the server

The implemented coordinator script realizes the main phases that are repeated during each run inside a loop and further collects the files containing the measurement data gathered on client side and on server side. The data evaluation based on that files after the test’s completion was conducted by specifically implemented R scripts.

C. Results and Limitations

The subsequent evaluation of the collected data within the unic²ast project enabled conclusions about the suitability of the DCCI system regarding its specific requirements, but are not subject of this paper. However, by benchmarking the real DCCI system according to the scenario defined above, it was possible to also draw conclusions about the suitability of the presented benchmark solution itself, which are explained below.

The test’s preparation and execution revealed some practical limitations of the presented approach, particularly with regard to the implementation in Bash and the usage of Qemu virtualization.

Firstly, the virtual clients cause a relatively high overhead in resource consumption. Regarding the amount of working memory, prior experience with lower numbers of clients already showed that each virtual client must be started with at least 16 MB RAM for the application to be stable, but in this case approximately 25 MB of host RAM are actually used per client, which speaks for a relatively high overhead due to the virtualization using Qemu. In order to carry out tests with the scenario described above with up to 1,000 clients, we therefore decided to use a host machine configuration equipped with 32 GB working memory, also taking into account some reserves to be used by the host’s operating system and still leaving about 5 to 6 GB of RAM free. Although the maximum number of clients was limited by the scenario, we aimed to further increase their number using the remaining free memory. But regarding the CPU usage, our tests showed this to be a limiting factor as well. Using a machine configuration with 8 CPU cores we could not run more than about 1,020 clients. We suspect that this is due to peaks in CPU usage during specific moments in the client lifecycle, e.g., during Zephyr kernel startup, configuration parameter provisioning, and temperature sensor simulation. However, the chosen system configuration proved to be sufficient for load tests with up to 1,000 virtualized clients. Ultimately, despite these observations, the

essential positive aspect of this approach should not be underestimated. The development of the test application brought valuable insights into the Zephyr operating system and its numerous interfaces. It can be considered close to the real software development for embedded solutions. This can also be particularly advantageous in the further course of the unic²ast project.

Secondly, while Bash scripting proved to be a suitable way to quickly adapt specific ideas into reality, it also clearly showed limitations. When the amount of source code grows over time during a software project, it is likely to run into specific issues regardless of the programming language used. Though simple modularization of Bash source is possible by separating them over several files, it lacks more sophisticated means to manage evolving dependencies. Furthermore, the absence of multi-dimensional arrays and the passing of data between different functions, both being essential elements in software programming, turned out to be challenges when using Bash. Ultimately, however, the solution could still be implemented, which is also due to the many helpful tools of the Linux ecosystem, e.g., `prips`.

V. CONCLUSION

The test setup described in the section above was successful overall and the listed requirements (see Section II) are demonstrably met by the presented approach:

- Qemu virtualization is able to run Zephyr OS based application.
- Qemu virtualization is used to run multiple Zephyr devices in parallel on the same host.
- The Lwm2m client in Zephyr serves as the basis for the development of a client application that enables a DTLS secured connection to the target system and can generate additional load on the target system through simulated Lwm2m resources.
- By using named pipes on the host, a bidirectional connection with the Zephyr OS based applications shell subsystem is used for the runtime parameterization of the clients and the interception of all outputs.
- The solutions presented are automated in the form of Bash script functions and can therefore be used in different test scenarios.

However, the test also revealed some limitations of the approach in terms of implementation difficulties due to Bash scripting and high resource consumption due to the overhead introduced by the virtualization mechanism.

Considering a possible further development of the presented benchmark setup, we suppose a replacement of the Bash based coordinator to be reasonable.

This would not only allow to use a advanced programming language like Java and its potentially large ecosystem of tools and frameworks. Moreover, it could let us implement the solution as an adapter for integration with existing tools specialized in load testing tool like JMeter. This allows us to leverage their potentials while not losing our ability to create realistic test scenarios based on real-world embedded software components. Over more scattering virtual client devices among multiple host systems could be realized. That will decrease the resource-based limits in the number of simulated devices.

Despite the high overhead in resource consumption, the use of emulated embedded devices with real application software has proven to be practical. The additional development of synthetic devices in network simulators or other specialized frameworks is not necessary. At the same time, the realized device behavior is quite close to the real world without the additional maintenance of a second code base. Furthermore, our approach tests the actual DCCI implementation and not a statistical simulation of it. Therefore, we expect test results which correspond to the real application.

On an overall view, our approach fulfils our formulated requirements and enables us to implement more complex load tests with minimum effort and the greatest possible accuracy.

REFERENCES

- [1] Statista Research Department. Internet of things - number of connected devices worldwide 2015-2025. [Online]. Available: <https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/> [retrieved: Dec., 2019]
- [2] unic²ast Projekt. [Online]. Available: <http://swt.informatik.uni-jena.de/Projekte/> [retrieved: Aug., 2020]
- [3] nsnam. ns-3 — a discrete-event network simulator for internet systems. [Online]. Available: <https://www.nsnam.org/> [retrieved: Feb., 2020]
- [4] OpenSim Ltd. OMNeT++ discrete event simulator. [Online]. Available: <https://omnetpp.org/> [retrieved: Jan., 2020]
- [5] Apache Software Foundation. Apache JMeter. [Online]. Available: <http://jmeter.apache.org/> [retrieved: Dec., 2019]
- [6] Gatling Corp. Gatling Open-Source Load Testing - For DevOps and CI/CD. [Online]. Available: <https://gatling.io/> [retrieved: Jan., 2020]
- [7] Open Mobile Alliance. Lightweight M2M (LWM2M) - OMA SpecWorks. [Online]. Available: <https://www.omaspecworks.org/what-is-oma-specworks/iot/lightweight-m2m-lwm2m/> [retrieved: Aug., 2019]
- [8] Zephyr Project. [Online]. Available: <https://www.zephyrproject.org/> [retrieved: Nov., 2019]
- [9] F. Bellard. QEMU. [Online]. Available: <https://www.qemu.org/> [retrieved: Jan., 2020]
- [10] Zephyr Project. Shell. [Online]. Available: <https://docs.zephyrproject.org/2.1.0/reference/shell/index.html> [retrieved: Apr., 2020]